

# Clusters and water flows: a novel approach to modal clustering through Morse theory

José E. Chacón\*

January 11, 2013

## Abstract

The problem of finding groups in data (cluster analysis) has been extensively studied by researchers from the fields of Statistics and Computer Science, among others. However, despite its popularity it is widely recognized that the investigation of some theoretical aspects of clustering has been relatively sparse. One of the main reasons for this lack of theoretical results is surely the fact that, unlike the situation with other statistical problems as regression or classification, for some of the cluster methodologies it is quite difficult to specify a population goal to which the data-based clustering algorithms should try to get close. This paper aims to provide some insight into the theoretical foundations of the usual nonparametric approach to clustering, which understands clusters as regions of high density, by presenting an explicit formulation for the ideal population clustering.

---

\*Departamento de Matemáticas, Universidad de Extremadura, E-06006 Badajoz, Spain. E-mail: [jechacon@unex.es](mailto:jechacon@unex.es)

# 1 Introduction

Clustering is one of the branches of Statistics with more research activity in the recent years. As noted by Meilă (2007), “clustering is a young domain of research, where rigorous methodology is still striving to emerge”. Indeed, some authors have recently expressed their concerns about the lack of theoretical or formal developments for clustering (von Luxburg and Ben-David, 2005; Ackerman and Ben-David, 2008; Zadeh and Ben-David, 2009). This paper aims to contribute to this regularization (or, say, rigorousization).

Stated in its most simple form, cluster analysis consists in “partitioning a data set into groups so that the points in one group are similar to each other and are as different as possible from the points in other groups” (Hand, Mannila and Smyth, 2001, p. 293). Posed as such, the problem does not even seem to have a statistical meaning.

In fact, some clustering techniques are solely based on the distances between the observations. Close observations are joined together to form a group, and extending the notion of inter-point distance to distance between groups, the resulting groups are gradually merged until all the initial observations are contained into a single group. This represents, of course, the notion of agglomerative *hierarchical clustering* (Izenman, 2008, Section 12.3). The graphical outcome depicting the successive agglomeration of data points up to a single group is the well-known dendrogram, and depending on the notion of inter-group distance used along the merging process the most common procedures of this type are known as single linkage, complete linkage or average linkage (see also Hastie, Tibshirani and Friedman, 2009, p. 523).

A first statistical flavour is noticed when dealing with those clustering methodologies that represent each cluster by a central point, such as the mean, the median or, more generally, a trimmed mean. This class of techniques is usually referred to as *partitioning methods*, and surely the most popular of its representatives is  $K$ -means (MacQueen, 1967). For a pre-specified number  $K$  of groups, these algorithms seek for  $K$  centers with the goal of optimizing a certain score function representing the quality of the clustering (Everitt *et al.*, 2011, Chapter 5).

A fully statistical notion of the problem is *distribution-based clustering*. This approach further assumes that the data set is a sample from a probability distribution. As with all the statistical procedures, there are parametric and nonparametric methodologies for distribution-based clustering. Surely the golden standard of parametric clustering is achieved through mixture modeling as clearly described in Fraley and Raftery (2002). It is assumed that the distribution generating the data is a mixture of simple parametric distributions, e.g. multivariate normal distributions, and each component of the mixture is associated to a different population cluster. Maximum likelihood is used to fit a mixture model and then each data point is assigned to the most likely component using the Bayes rule.

Eventually, by using mixture distribution components defined through Mahalanobis distances, some partitioning methods and mixture-model clustering can be regarded to have the same motivation, and under this point of view some partitioning techniques have evolved to reach a high level of maturity, leading to very sophisticated methods. See, for instance, the recent review of García-Escudero *et al.* (2010) and references therein.

In parametric distribution-based clustering the notion of population cluster is clearly defined through the components of the mixture model. If the true underlying distribution belongs indeed to the postulated mixture model, this identification is certainly quite adequate. However, often mixture models are used in a “nonparametric way”, on the basis that they usually represent a dense subclass of the set of all continuous distributions and, therefore, they are useful to approximate arbitrarily well any continuous distribution. In that case, the identification of components with clusters may not coincide with the natural intuition. For instance, several normal components may be needed to model accurately a distribution with a heavier-than-normal tail, yet the resulting normal mixture might as well be a unimodal density, suggesting the existence of a single group, though with a non-normal distribution shape. The fact that mixture components and clusters are not necessarily the same has motivated several variations of mixture-model clustering to amend this deficiency; see, for example, Biernacki, Celeux and Govaert (2000), Baudry *et al.* (2010) or Hennig (2010).

The previous example is helpful to understand the motivation of the nonparametric methodology, which is based on identifying clusters as regions of high density separated from each other by regions of lower density. In this sense, population clusters are naturally associated with the modes (i.e., local maxima) of the probability density function, and this nonparametric approach is denominated mode-based clustering or *modal clustering* (Li, Ray and Lindsay, 2007). Precisely, each cluster is usually understood as the “domain of attraction” of a mode (Stuetzle, 2003).

The concept of domain of attraction is not that simple to specify, and providing a precise definition is one of the goals of this paper that will be explored in the next section. It is perhaps this imprecise notion of population cluster in the nonparametric case what does not fully convince many researchers and statisticians and lead them to adopt one of the alternative previously expounded techniques (parametric clustering, partitioning methods or hierarchical algorithms).

The first attempt to make the goal of modal clustering precise is through the notion of level sets (Hartigan, 1975). If the distribution of the data has a density  $f$ , given  $\lambda \geq 0$  the  $\lambda$ -level set of  $f$  is defined as  $L(\lambda) = \{\mathbf{x}: f(\mathbf{x}) \geq \lambda\}$ . Then, population  $\lambda$ -clusters are defined as the connected components of  $L(\lambda)$ , a definition that clearly captures the notion of groups having a high density. An extensive account of the usefulness of level sets in applications is given in Mason and Polonik (2009).

The main drawback of this approach is the fact that the notion of population cluster




Figure 1: Univariate trimodal density for which it is not possible to capture its whole cluster structure using a level set analysis based on a single level.

depends on the level  $\lambda$ , as recognized by Stuetzle (2003). However, other authors like Cuevas, Febrero and Fraiman (2001) or Cadre, Pelletier and Pudlo (2011), affirm that the choice of  $\lambda$  is only a matter of resolution level of the analysis. Indeed, the level sets can be re-parameterized in terms of their probability content, as in Cadre (2006) or Azzalini and Torelli (2007), for ease of interpretability. This is usually done by finding, for any  $p \in (0, 1)$ , the threshold  $\lambda_p$  defined as the largest value of  $\lambda$  such that the probability of  $L(\lambda)$  is greater than  $p$ , and setting  $R(p) = L(\lambda_p)$  so that the level set  $R(p)$  has probability  $p$  as long as the border of  $R(p)$  has null probability. Then, the chosen value of  $p$  implies that a proportion  $\alpha = 1 - p$  of the data is left out of the clustering mechanism, and therefore the level  $\lambda$  (or  $p$ , or  $\alpha$ ) can be interpreted as a trimming parameter motivated by robustness considerations. In fact, the parametrization of the level sets in terms of the probability  $\alpha$  of the set of non-classified observations is also commonly used; see Hyndman (1996); Baíllo, Cuesta-Albertos and Cuevas (2001) or Samworth and Wand (2010).

Nevertheless, it is easy to think of many examples in which it is impossible to observe the whole cluster structure on the basis of a single level  $\lambda$ . Essentially as in Rinaldo *et al.* (2012, p. 906), Figure 1 shows a simple univariate example of this phenomenon: three different groups are visually identifiable, yet none of the level sets of the density has three connected components. To amend this, the usual recommendation is to analyze the cluster structure for several values of the level  $\lambda$  (or the coverage probability  $p$ ). Graphical tools oriented to this goal are the cluster tree (Stuetzle, 2003), or the mode function (Azzalini and Torelli, 2007). Both graphics are useful to show how the clusters emerge as a function of  $\lambda$  or  $p$ , respectively. A detailed description of the cluster identification methodology through the cluster tree is given in Section 2 below.

Level sets also arise as the maximizers of excess mass functionals. For a given value of  $\lambda \geq 0$ , the excess mass functional is defined as  $E_\lambda(A) = P(A) - \lambda\mu(A)$  for any measurable set  $A$ , where  $\mu$  denotes the Lebesgue measure. Hartigan (1987) and Müller and Sawitzki (1987, 1991) independently introduced this functional and noticed that it reaches its maximum value precisely at  $L(\lambda)$ . A nice feature of the excess mass functional is that it admits

a fairly simple empirical estimator obtained by replacing  $P$  for the empirical probability measure of the sample. However, the maximizer of this empirical criterion does not provide a useful estimate of  $L(\lambda)$  because it turns out to be just the set of all the data points. Nevertheless, by restricting the empirical functional to a suitable class of measurable sets its maximizer can be shown to be a sensible estimate of  $L(\lambda)$  within such a class. This estimation procedure is studied in detail in Polonik (1995) making use of empirical processes theory.

On the other hand, the estimation of density level sets using smoothing techniques has its roots in the papers of Cuevas and Fraiman (1997); Tsybakov (1997) or Walther (1997), and related data-driven cluster methodologies based on the notion of  $\lambda$ -clusters are proposed in Cuevas, Febrero and Fraiman (2001); Stuetzle (2003); Azzalini and Torelli (2007); Nugent and Stuetzle (2010); Stuetzle and Nugent (2010); Rinaldo and Wasserman (2010) or Rinaldo *et al.* (2012), among others. In all these papers a plug-in approach is adopted, in the sense that a population level set is estimated by the corresponding level set of some density estimator. It is to be noted that for most of the plug-in estimators, and also for the excess mass estimators, the computation of empirical level sets is indeed a difficult issue, and many of the aforementioned references also present intricate algorithms to approximate the corresponding level set estimators. Relatedly, in Machine Learning literature surely the most popular algorithms for density-based clustering include DBSCAN (Easter *et al.*, 1996), based on a single level, and OPTICS (Ankerst *et al.*, 1999), which is a generalization of the former which allows to compute the cluster tree induced by a range of levels.

## 2 A precise approach to the goal of modal clustering

Certainly, the very concept of modal cluster may admit some well-supported criticism, as illustrated in the thorough reflections of Stuetzle (2003, Section 6) or Hennig (2010), and we agree with Izenman (2008, p. 408) that “there is no universally accepted definition of exactly what constitutes a cluster”. Some of the most common definitions have been surveyed in the previous section. However, it is not the goal of this paper to contribute to the discussion of the different notions of cluster, but to introduce a precise definition of population modal cluster which does not rely on the concept of level set. This aims to provide an answer to Question 1 in von Luxburg and Ben-David (2005): “How does a desirable clustering look like if we have complete knowledge about our data generating process?”

Since the empirical formulation of the clustering task consists of partitioning a data set into homogeneous groups, a population clustering  $\mathcal{C}$  of a probability distribution  $P$  on  $\mathbb{R}^d$  should be understood as a partition of  $\mathbb{R}^d$  into mutually disjoint measurable components  $C_1, \dots, C_k$ , each with positive probability content (Ben-David, von Luxburg and Pál, 2006). The components of such a partition are called population clusters. Thus, two population

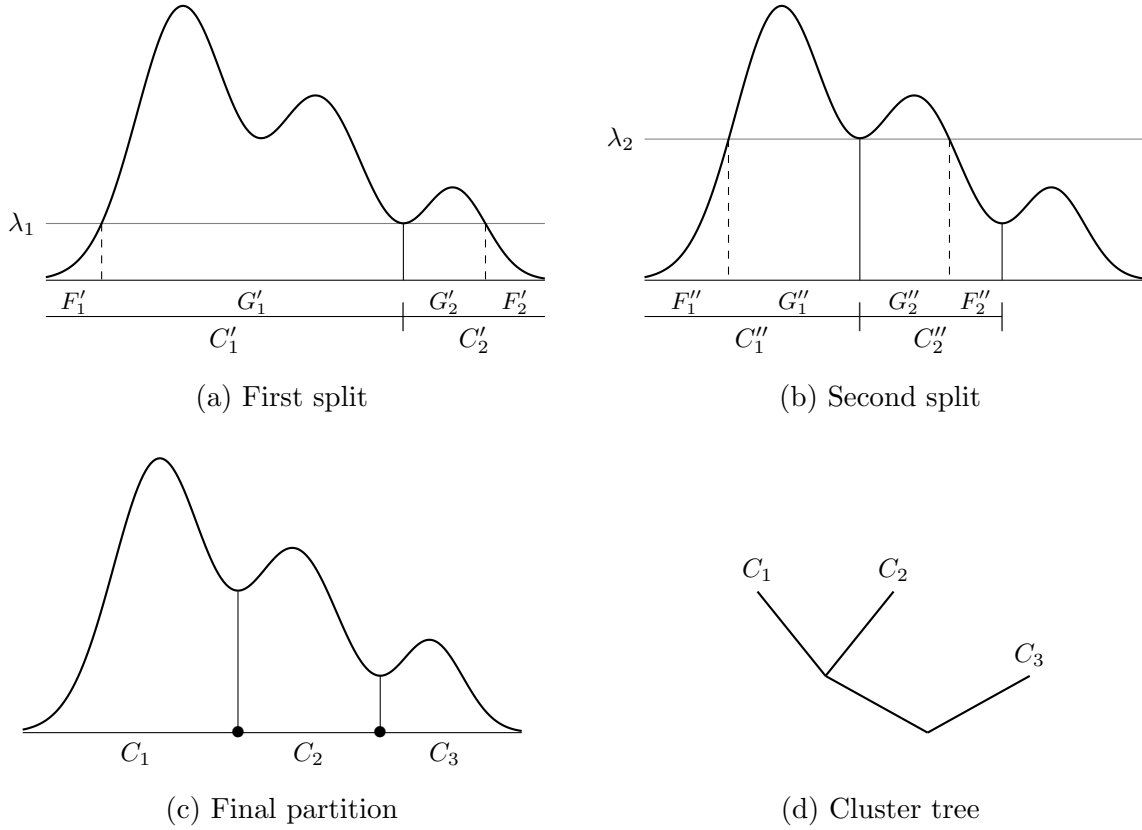


Figure 2: Identification of clusters for the trimodal density example using the cluster tree.

clusterings  $\mathcal{C}$  and  $\mathcal{D}$  are identified to be the same if they have the same number of clusters and, up to a permutation of the cluster labels, every cluster of  $\mathcal{C}$  and its most similar match in  $\mathcal{D}$  differ in a null-probability set (more details on this are elaborated in Section 3).

As anticipated in the previous section, a basic principle about the concept of ideal population clustering is that it should reflect the notion of a partition into regions of high density separated from each other by regions of lower density, leading to the idea of population cluster as the domain of attraction of a density mode. We begin our way to formalize this concept by showing some illustrative examples in one and two dimensions.

In the one-dimensional case, it seems clear from Figure 2 how this can be achieved. To begin with, the level set methodology identifies the three clusters in the density depicted in Figure 1 by computing the cluster tree as described clearly in Nugent and Stuetzle (2010): starting from the 0-level set, which corresponds to the whole real line in this example and consists of a single connected component,  $\lambda$  is increased until it reaches  $\lambda_1$ , where two components for the  $\lambda_1$ -level set are found,  $G'_1$  and  $G'_2$ , resulting in the cluster tree splitting into two different branches (see Fig. 2, panel (a)). These two components  $G'_1$  and  $G'_2$  are usually called cluster cores. They do not constitute a partition of  $\mathbb{R}$ , but the remaining parts  $F'_1$  and  $F'_2$ , referred to as fluff in Nugent and Stuetzle (2010), can be assigned to either




Figure 3: Bidimensional example, with two groups clearly identifiable at an intuitive level.

the left or the right branch depending on whichever of them is closer. Thus, at level  $\lambda_1$  the partition  $\mathbb{R} = C'_1 \cup C'_2$  is obtained. The point dividing the line into these two components can be arbitrarily assigned to either of them; this assignment makes no difference, because it leads to equivalent clusterings.

At level  $\lambda_2$  the left branch is further divided into two branches (see panel (b) of Fig. 2). Again, the two cluster core components  $G''_1$  and  $G''_2$  do not form a partition of the set  $C'_1$  associated with the previous node of the tree, but it is clear how the fluff  $F''_1$  and  $F''_2$  can be assigned to form a partition  $C''_1 \cup C''_2$  of  $C'_1$ . Since no further splitting of the cluster tree is observed as  $\lambda$  increases, the final partition of  $\mathbb{R}$  is  $C''_1 \cup C''_2 \cup C'_2$ , renamed to  $C_1 \cup C_2 \cup C_3$  in panel (c) of Figure 2.

It is immediate to observe that the levels at which a connected component breaks into two different ones correspond precisely to local minima of the density function, so an equivalent formulation consists of defining population clusters as the connected components of  $\mathbb{R}$  minus the points where a local minimum is attained (the solid circles in panel (c) of Figure 2). Notice that this definition does not involve the computation of level sets for a range of levels, nor their cores and fluff, so that it constitutes a more straightforward approach to the same concept in the unidimensional setup.

To get an idea of how to generalize the previous approach to higher dimensions, consider the following extremely simple bidimensional example: an equal-proportion mixture of two normal distributions, each with identity variance matrix and centered at  $\mu_1 = (-\frac{3}{2}, 0)$  and  $\mu_2 = -\mu_1$ , respectively. At an intuitive level, it is clear from Figure 3 that the most natural border to separate the two visible groups is the black line. The problem is: what is exactly that line? Is it identifiable in terms of the features of the density function in a precise, unequivocal way? A nice way to answer to these questions is by means of Morse theory.




Figure 4: The three possible configurations around a critical point of a Morse function in the bidimensional case.

Morse theory is a branch of Differential Geometry that provides tools for analyzing the topology of a manifold  $M \subseteq \mathbb{R}^d$  by studying the critical points of a differentiable function  $f: M \rightarrow \mathbb{R}$ . A classical reference book on this subject is Milnor (1963) and enjoyable introductions to the topic can be found in Matsumoto (2002) and Jost (2011, Chapter 7). A useful application of Morse theory is for terrain analysis, as nicely developed in Vitalli (2010). In terrain analysis, a mountain range can be regarded as the graph of a function  $f: M \rightarrow \mathbb{R}$ , representing the terrain elevation, over a terrain  $M \subseteq \mathbb{R}^2$ , just as in the left graphic of Figure 3. The goal of terrain analysis is to provide a partition of  $M$  through watersheds indicating the different regions, or catchment basins, where water flows under the effect of gravity.

The fundamentals of Morse theory can be extremely summarized as follows. A smooth function  $f: M \rightarrow \mathbb{R}$  is called a *Morse function* if all its critical points are non-degenerate. Here the critical points of  $f$  are understood as those  $\mathbf{x}_0 \in M$  for which the gradient  $Df(\mathbf{x}_0)$  is null, and non-degeneracy means that the determinant of the Hessian matrix  $Hf(\mathbf{x}_0)$  is not zero. For such points the *Morse index*  $m(\mathbf{x}_0)$  is defined as the number of negative eigenvalues of  $Hf(\mathbf{x}_0)$ .

Morse functions can be expressed in a fairly simple form in a neighborhood of a critical point  $\mathbf{x}_0$ , as the result known as Morse lemma shows that it is possible to find local coordinates  $x_1, \dots, x_n$  such that  $f$  can be written as  $\pm x_1 \pm \dots \pm x_d + c$  around  $\mathbf{x}_0$ , where the number of minus signs in the previous expression is precisely  $m(\mathbf{x}_0)$ . For example, for  $d = 2$  the three possible configurations for a critical point are shown in Figure 4, corresponding to a local minimum, a saddle point and a local maximum (from left to right), with Morse indexes 0, 1 and 2, respectively.

The decomposition of  $M$  suggested by Morse theory is made in terms of the unstable and/or stable manifolds of the critical points of  $f$  as explained next. Consider the initial value problem defined by the minus gradient vector of  $f$ . For a given value of  $\mathbf{x} \in M$  at



time  $t = 0$ , the integral curve  $\nu_x: \mathbb{R} \rightarrow M$  of such initial value problem is the one satisfying

$$\nu'_x(t) = -Df(\nu_x(t)), \quad \nu_x(0) = x$$

and the set of all these integral curves is usually referred to as the negative gradient flow. Since the minus gradient vector defines the direction of steepest descent of  $f$ , these curves (or properly speaking, their images through  $f$ ) represent the trajectories of the water flow subject to gravity.

With respect to the negative gradient flow, the *unstable manifold* of a critical point  $x_0$  is defined as the set of points whose integral curve starts at  $x_0$ , that is

$$W_-^u(x_0) = \{x \in M: \lim_{t \rightarrow -\infty} \nu_x(t) = x_0\}.$$

Analogously, the stable manifold of  $x_0$  is the set of points whose integral curve finishes at  $x_0$ , that is,  $W_-^s(x_0) = \{x \in M: \lim_{t \rightarrow +\infty} \nu_x(t) = x_0\}$ . It was firstly noted by Thom (1949) that the set of unstable manifolds corresponding to all the critical points of  $f$  provides a partition of  $M$  (the same is true for the stable manifolds). Furthermore, the unstable manifold  $W_-^u(x_0)$  has dimension  $m(x_0)$ .

The main contribution of this section is to define the population clusters of a density  $f$  as the unstable manifolds of the negative gradient flow corresponding to local maxima of  $f$ . Or in a more prosaic way, in terms of water flows, a cluster is just the region of the terrain that would be flooded by a fountain emanating from a peak of the mountain range.

Although this is an admittedly cumbersome definition, going back to Figure 3 it is clear that it just describes the notion that we were looking for. The critical point  $x_0 = (0, 0)$  is a saddle point, thus having Morse index 1, and the black line is precisely its associated unstable manifold,  $W_-^u(x_0) = \{0\} \times \mathbb{R}$ , which is a manifold of dimension 1. The remaining two critical points are local maxima, and their respective unstable manifolds are  $W_-^u(x_1) = (-\infty, 0) \times \mathbb{R}$  and  $W_-^u(x_2) = (0, \infty) \times \mathbb{R}$ , manifolds of dimension 2 so that we can partition  $\mathbb{R}^2 = W_-^u(x_0) \cup W_-^u(x_1) \cup W_-^u(x_2)$ , showing  $W_-^u(x_1)$  and  $W_-^u(x_2)$  as two population clusters separated by the border  $W_-^u(x_0)$ , which is a null-probability set.

Notice that this definition also applies to the previous univariate example in Figure 2: the clusters  $C_1$ ,  $C_2$  and  $C_3$  are just the unstable manifolds of the three local maxima (they are manifolds of dimension 1), and for the two local minima their unstable manifolds have dimension 0, so they include only the respective points of minima.

Moreover, if we focus on the gradient flow, instead of the negative gradient flow, then its integral curves satisfy

$$\gamma'_x(t) = Df(\gamma_x(t)), \quad \gamma_x(0) = x,$$

the unstable manifold for the negative gradient flow becomes the stable manifold for the gradient flow and viceversa. Therefore, we could equivalently define the cluster associated to a mode  $x_0$  of the density as its stable manifold with respect to the gradient flow, that

is,  $W_+^s(\mathbf{x}_0) = \{\mathbf{x} \in M : \lim_{t \rightarrow \infty} \gamma_{\mathbf{x}}(t) = \mathbf{x}_0\}$ . This is a precise formulation of the notion of domain of attraction of the mode  $\mathbf{x}_0$ , since  $W_+^s(\mathbf{x}_0)$  represents the set of all the points that climb to  $\mathbf{x}_0$  when they follow the steepest ascent path defined by the gradient direction.

### 3 Comparing population clusterings

Once the population goal of the clustering problem has been clearly identified as the partition of  $\mathbb{R}^d$  given in terms of the stable manifolds of the density modes with respect to the gradient flow, it is natural to wonder how this population clustering can be estimated from the data, how the error between a data-based clustering and the population clustering can be measured and to set up a notion of consistency. Those are the goals for this section.

At this point it is worth distinguishing between two different, although closely related, concepts. When the density  $f$  is unknown, and a sample  $\mathbf{X}_1, \dots, \mathbf{X}_n$  drawn from  $f$  is given, any procedure to obtain a data-based partition of  $\mathbb{R}^d$  will be called a *data-based clustering*. Notice, however, that when data are available the goal of most clustering procedures is just to partition the data set, and not the whole space  $\mathbb{R}^d$ . This will be referred to henceforth as a *clustering of the data*. Naturally, the ideal population clustering or any data-based clustering immediately result in a clustering of the data, by assigning the same group to data points belonging to the same component of the given partition of  $\mathbb{R}^d$ .

A good deal of measures to evaluate the distance between two clusterings of the data have been proposed in the literature. The work of Meilă (2007) provides both a comprehensive survey of the most used existing measures as well as a deep technical study of their main properties, and for instance Arabie and Boorman (1973) or Day (1981) provide further alternatives. But apart from some exceptions, as the Hamming distance considered in Ben-David, von Luxburg and Pál (2006), it should be stressed that all these proposals concern partitions of a finite set. However, since the ideal population goal of clustering introduced here is defined as a partition of  $\mathbb{R}^d$ , it is necessary to introduce distances between two partitions of  $\mathbb{R}^d$  to evaluate the differences between two clusterings.

#### 3.1 A distance in measure between clusterings

Let  $\mathcal{C}$  and  $\mathcal{D}$  be two clusterings of a probability distribution  $P$ , and assume for the moment that both have the same number of clusters, say  $\mathcal{C} = \{C_1, \dots, C_r\}$  and  $\mathcal{D} = \{D_1, \dots, D_r\}$ . The first step to introduce a distance between  $\mathcal{C}$  and  $\mathcal{D}$  is to consider a distance between sets. Surely the two distances between sets most used in practice are the Hausdorff distance and the distance in measure; see Cuevas and Fraiman (2010). The Hausdorff distance is specially useful when dealing with compact sets (it defines a metric in the space of all compact sets of a metric space), it tries to capture the notion of physical proximity between two sets (Rodríguez-Casal, 2003). In contrast, the distance in measure between two sets  $C$

and  $D$  refers to the content of their symmetric difference  $C \triangle D = (C \cap D^c) \cup (C^c \cap D)$ , and it defines a metric on the set of all measurable subsets of a measure space, once two sets differing in a null-measure set are identified to be the same.

Although we will return to the Hausdorff distance later, our first approach to the notion of distance between  $\mathcal{C}$  and  $\mathcal{D}$  relies primarily on the concept of distance in measure, and the measure involved is precisely the probability measure  $P$ . From a practical point of view, it does not seem so important that the clusters of a data-based partition get physically close to those of the ideal clustering. Instead, it is desirable that the points that are incorrectly assigned do not represent a very significant portion of the distribution. Thus, it seems natural to affirm that two clusters  $C \in \mathcal{C}$  and  $D \in \mathcal{D}$  (resulting from different clusterings) are close when  $P(C \triangle D)$  is low.

Therefore, for two clusterings  $\mathcal{C}$  and  $\mathcal{D}$  with the same number of clusters, a sensible notion of distance is obtained by adding up the contributions of the pairwise distances between their components, once they have been relabelled so that every cluster in  $\mathcal{C}$  is compared with its most similar counterpart in  $\mathcal{D}$ . In mathematical terms, the distance between  $\mathcal{C}$  and  $\mathcal{D}$  can be measured by

$$d_1(\mathcal{C}, \mathcal{D}) = \min_{\sigma \in \mathcal{P}_r} \sum_{i=1}^r P(C_i \triangle D_{\sigma(i)}), \quad (1)$$

where  $\mathcal{P}_r$  denotes the set of permutations of  $\{1, 2, \dots, r\}$ .

It can be shown that  $d_1$  defines a metric in the space of all the partitions with the same number of components, once two such partitions are identified to be the same if they differ only in a relabelling of their components. Moreover, the minimization problem in (1) is usually known as the *linear sum assignment problem* in the literature of Combinatorial Optimization, and it represents a particular case of the well-known Monge-Kantorovich transportation problem. A comprehensive treatment of assignment problems can be found in Burkard, Dell'Amico and Martello (2009).

If a partition is understood as a vector in the product space of measurable sets, with the components as its coordinates, then  $d_1$  resembles the  $L_1$  product distance, only adapted to take into account the possibility of relabelling the components. This seems a logical choice given the additive nature of measures, as it adds up the contribution of each distance between the partition components as described before. However, it would be equally possible to consider any other  $L_p$  distance, leading to define

$$d_p(\mathcal{C}, \mathcal{D}) = \min_{\sigma \in \mathcal{P}_r} \left\{ \sum_{i=1}^r P(C_i \triangle D_{\sigma(i)})^p \right\}^{1/p}$$

for  $p \geq 1$  and also  $d_\infty(\mathcal{C}, \mathcal{D}) = \min_{\sigma \in \mathcal{P}_r} \max\{P(C_i \triangle D_{\sigma(i)}): i = 1, \dots, r\}$ . The minimization problem defining  $d_\infty$  is also well known, under the name of the linear bottleneck assignment problem, and its objective function is usually employed if the interest is to minimize the




Figure 5: Two partitions of the unit square that do not differ much if  $A_3$  has low probability.

latest completion time in parallel computing (see Burkard, Dell’Amico and Martello, 2009, Section 6.2). In the context of clustering, surely the  $d_1$  distance seems the most natural choice among all the  $d_p$  possibilities, due to its clear interpretation.

Nevertheless, the definition of the  $d_1$  distance involves some kind of redundancy, due to the fact that  $\mathcal{C}$  and  $\mathcal{D}$  are partitions of  $\mathbb{R}^d$ , because the two disjoint sets that form every symmetric difference in fact appear twice in each of the sums in (1). More precisely, taking into account that  $P(C \triangle D) = P(C) + P(D) - 2P(C \cap D)$ , it follows that for every  $\sigma \in \mathcal{P}_r$

$$\sum_{i=1}^r P(C_i \triangle D_{\sigma(i)}) = 2 - 2 \sum_{i=1}^r P(C_i \cap D_{\sigma(i)}). \quad (2)$$

To avoid this redundancy, the eventual suggestion to measure the distance between  $\mathcal{C}$  and  $\mathcal{D}$ , based on the set distance in  $P$ -measure, is  $d_P(\mathcal{C}, \mathcal{D}) = \frac{1}{2}d_1(\mathcal{C}, \mathcal{D})$ .

If the partitions  $\mathcal{C}$  and  $\mathcal{D}$  do not have the same number of clusters, then as many empty set components as needed are added so that both partitions include the same number of components, as in Charon *et al.* (2006), and the distance between the extended partitions is computed as before. Explicitly, if  $\mathcal{C} = \{C_1, \dots, C_r\}$  and  $\mathcal{D} = \{D_1, \dots, D_s\}$  with  $r < s$  then, writing  $C_i = \emptyset$  for  $i = r+1, \dots, s$ , we set

$$d_P(\mathcal{C}, \mathcal{D}) = \frac{1}{2} \min_{\sigma \in \mathcal{P}_s} \sum_{i=1}^s P(C_i \triangle D_{\sigma(i)}) = \frac{1}{2} \min_{\sigma \in \mathcal{P}_s} \left\{ \sum_{i=1}^r P(C_i \triangle D_{\sigma(i)}) + \sum_{i=r+1}^s P(D_{\sigma(i)}) \right\}.$$

Thus, the term  $\sum_{i=r+1}^s P(D_{\sigma(i)})$  can be interpreted as a penalization for unmatched probability mass.

The idea is that two partitions as those shown in Figure 5 do not differ much if  $A_3$  has low probability, even if they do not have the same number of clusters. For the partitions in Figure 5, denote  $\mathcal{C} = \{C_1, C_2\}$  and  $\mathcal{D} = \{D_1, D_2, D_3\}$  with  $C_1 = D_1 = A_1$ ,  $C_2 = A_2 \cup A_3$ ,  $D_2 = A_2$ ,  $D_3 = A_3$ , and assume that  $P(A_1) = 0.5$ ,  $P(A_2) = 0.45$  and  $P(A_3) = 0.05$ . Then, it can be shown that  $d_P(\mathcal{C}, \mathcal{D}) = 0.05$ . In (1) every cluster of  $\mathcal{C}$  is matched to some cluster in  $\mathcal{D}$ , depending on the permutation for which the minimum is achieved. When  $\mathcal{C}$

has less clusters than  $\mathcal{D}$ , some of the components of  $\mathcal{D}$  will be matched with the empty set, indicating that they do not have an obvious match in  $\mathcal{C}$ , or that they are unimportant. In the previous example, the minimum is achieved when  $C_1$  is matched with  $D_1$ ,  $C_2$  with  $D_2$  and  $D_3$  is matched with the empty set.

It is interesting to note that  $d_P(\mathcal{C}, \mathcal{D})$  can be estimated in a natural way by replacing  $P$  with the empirical measure based on the data  $\mathbf{X}_1, \dots, \mathbf{X}_n$ , leading to

$$\widehat{d}_P(\mathcal{C}, \mathcal{D}) = \frac{1}{2n} \min_{\sigma \in \mathcal{P}_s} \left\{ \sum_{i=1}^r \sum_{j=1}^n I_{C_i \Delta D_{\sigma(i)}}(\mathbf{X}_j) + \sum_{i=r+1}^s \sum_{j=1}^n I_{D_{\sigma(i)}}(\mathbf{X}_j) \right\}$$

where  $I_A$  denotes the indicator function of the set  $A$ . When  $r = s$ , it follows from (2) that an alternative expression for  $d_P(\mathcal{C}, \mathcal{D})$  is

$$d_P(\mathcal{C}, \mathcal{D}) = 1 - \max_{\sigma \in \mathcal{P}_r} \sum_{i=1}^r P(C_i \cap D_{\sigma(i)})$$

and, therefore, its sample analogue

$$\widehat{d}_P(\mathcal{C}, \mathcal{D}) = 1 - \frac{1}{n} \max_{\sigma \in \mathcal{P}_r} \sum_{i=1}^r \sum_{j=1}^n I_{C_i \cap D_{\sigma(i)}}(\mathbf{X}_j),$$

coincides with the so-called classification distance between two clusterings of the data, whose properties are explored in Meilă (2005, 2007, 2012). For  $r < s$ , however,  $\widehat{d}_P$  differs from the classification distance (which does not include the penalty term), but it corresponds exactly with the transfer distance, studied in detail in Charon *et al.* (2006) (see also Dencéud, 2008). Extending the properties of the transfer distance to its population counterpart suggests an interpretation of  $d_P(\mathcal{C}, \mathcal{D})$  as the minimal probability mass that needs to be moved to transform the partition  $\mathcal{C}$  into  $\mathcal{D}$ , hence the connection with the optimal transportation problem.

### 3.2 A distance between sets of sets

An alternative notion of distance between two clusterings based on the Hausdorff metric has been kindly suggested by Professor Antonio Cuevas, noting that precisely this distance was used in Pollard (1981) to measure the discrepancy between the set of sample  $K$ -means and the set of population  $K$ -means. If  $(X, \rho)$  is a metric space and  $A, B \subseteq X$  are two non-empty subsets of  $X$ , the Hausdorff distance between  $A$  and  $B$  is defined as

$$d_H(A, B) = \max \left\{ \sup_{a \in A} \inf_{b \in B} \rho(a, b), \sup_{b \in B} \inf_{a \in A} \rho(a, b) \right\}$$

or, equivalently, as

$$d_H(A, B) = \inf \{ \varepsilon > 0 : A \subseteq B^\varepsilon \text{ and } B \subseteq A^\varepsilon \},$$

where  $A^\varepsilon = \bigcup_{a \in A} \{x \in X : \rho(x, a) \leq \varepsilon\}$ , and  $B^\varepsilon$  is defined analogously.

In the context of clustering, take  $X$  to be the metric space consisting of all the sets of  $\mathbb{R}^d$  equipped with the distance  $\rho(C, D) = P(C \triangle D)$ , once two sets with  $P$ -null symmetric difference have been identified to be the same. Then any two clusterings  $\mathcal{C} = \{C_1, \dots, C_r\}$  and  $\mathcal{D} = \{D_1, \dots, D_s\}$  can be viewed as (finite) subsets of  $X$ , and therefore the Hausdorff distance between  $\mathcal{C}$  and  $\mathcal{D}$  is defined as

$$\begin{aligned} d_H(\mathcal{C}, \mathcal{D}) &= \max \left\{ \max_{i=1, \dots, r} \min_{j=1, \dots, s} P(C_i \triangle D_j), \max_{j=1, \dots, s} \min_{i=1, \dots, r} P(C_i \triangle D_j) \right\} \\ &= \inf \{ \varepsilon > 0 : \mathcal{C} \subseteq \mathcal{D}^\varepsilon \text{ and } \mathcal{D} \subseteq \mathcal{C}^\varepsilon \}. \end{aligned}$$

To express it in words,  $d_H(\mathcal{C}, \mathcal{D}) \leq \varepsilon$  whenever for every  $C_i \in \mathcal{C}$  there is some  $D_j \in \mathcal{D}$  such that  $P(C_i \triangle D_j) \leq \varepsilon$  and viceversa.

The Hausdorff distance can be regarded as a uniform distance between sets. It is not hard to show, using standard techniques from the Theory of Normed Spaces, that when  $r = s$  we have

$$d_H(\mathcal{C}, \mathcal{D}) \leq 2 d_P(\mathcal{C}, \mathcal{D}) \leq r d_H(\mathcal{C}, \mathcal{D}).$$

However, when  $r < s$  the distance  $d_H$  can be more demanding than  $d_P$ , meaning that both partitions have to be really close so that their Hausdorff distance results in a small value. For instance, it can be checked that for the two clusterings of the previous example, shown in Figure 5, the Hausdorff distance between them is  $d_H(\mathcal{C}, \mathcal{D}) = 0.45$ , mainly due to the fact that  $C_2$  and  $D_3$  are far from each other, since  $P(C_2 \triangle D_3) = P(A_2) = 0.45$ . A clear picture of the difference between  $d_H$  and  $d_P$  is obtained by arranging all the component-wise distances  $P(C_i \triangle D_j)$  into an  $r \times s$  matrix. Then, the Hausdorff distance is obtained by computing all the row-wise and column-wise minima and taking the maximum of all of them. In contrast, for the distance in  $P$ -measure the first step when  $r < s$  is to add  $s - r$  row copies of the vector  $(P(D_1), \dots, P(D_s))$  to the matrix of component-wise distances, and then compute the distance in  $P$ -measure as half the minimum possible sum obtained by adding up a different element in each row. As a further difference, note that the Hausdorff distance does not involve a matching problem; instead, this distance is solely determined by the two components that are furthest from each other.

Obviously, a sample analogue is also obtained in this case by replacing  $P$  for the empirical probability measure, leading to

$$\hat{d}_H(\mathcal{C}, \mathcal{D}) = \frac{1}{n} \max \left\{ \max_{i=1, \dots, r} \min_{j=1, \dots, s} \sum_{k=1}^n I_{C_i \triangle D_j}(\mathbf{X}_k), \max_{j=1, \dots, s} \min_{i=1, \dots, r} \sum_{k=1}^n I_{C_i \triangle D_j}(\mathbf{X}_k) \right\},$$

which seems not to have been considered previously as a distance between two clusterings of the data.




Figure 6: Two density functions that are not close but induce exactly the same clustering.

### 3.3 Consistency of data-based clusterings

As indicated above, a data-based clustering is understood as any procedure that induces a partition  $\hat{\mathcal{C}}_n$  of  $\mathbb{R}^d$  based on the information obtained from a sample  $\mathbf{X}_1, \dots, \mathbf{X}_n$  from a probability distribution  $P$ . Once we have a precise definition for the ideal population clustering goal  $\mathcal{C}$ , a data-based clustering  $\hat{\mathcal{C}}_n$  is said to be consistent if it gets closer to  $\mathcal{C}$  as the sample size increases. Formally, if  $d(\hat{\mathcal{C}}_n, \mathcal{C}) \rightarrow 0$  as  $n \rightarrow \infty$  for some of the modes of stochastic convergence (in probability, almost surely, etc), where  $d$  represents one of the distances between clusterings defined above, or any other sensible alternative.

A plug-in strategy to obtain data-based clusterings naturally arises when the density  $f$  is replaced with a smooth estimator  $\hat{f}$ , and so the data-based clusters are defined as the domains of attraction of the modes of  $\hat{f}$ . Obvious candidates for the role of  $\hat{f}$  include kernel density estimators for the nonparametric setup or mixture model density estimators in a parametric context.

This is a very simple approach, that involves to some extent estimating the density function to solve the clustering problem (unsupervised learning), and according to von Luxburg (2004, p. 21) it may not be a good idea because density estimation is a very difficult problem, especially in high dimensions. A similar situation is found in the study of classification (supervised learning), where the optimal classifier, the Bayes rule, depends on the regression function of the random labels over the covariates. Here, even if classification can be proved to be a problem easier than regression, nevertheless regression-based algorithms for classification play an important role in the development of supervised learning theory (see Devroye, Györfi and Lugosi, 1996, Chapter 6).

Along the same lines, Figure 6 illustrates why we should not completely discard density estimation as an intermediate step for clustering. Figure 6 shows a typical situation where the solid line is the true density and the dashed line is a kernel density estimator, since an expansion of its pointwise bias shows that on average the kernel estimator underestimates the maxima and overestimates the minima (Wand and Jones, 1995, p. 21). But even if the two density functions are not really close in any global sense, they produce exactly the same clusterings of  $\mathbb{R}$ , so this also seems to suggest that clustering should be easier than density

estimation. Studying consistency and error rates for data-based clusterings obtained from density estimators surely represents an interesting open problem.

## 4 Practical algorithms

Apart from the theoretical considerations in the previous sections, for practical purposes it is necessary to provide algorithms that help to compute population or data-based clusterings. As noted in Section 1, even if the density  $f$  is fully known this may become an onerous task for some methodologies, as for instance those based on level sets, for which smart implementations are needed to compute the connected components of the level sets or the cluster tree (Cuevas, Febrero and Fraiman, 2001; Stuetzle and Nugent, 2010).

With the concept of cluster as the domain of attraction of the density modes, the natural approach to compute the partition induced by a density is to use the so-called mode seeking algorithms. Surely the most popular algorithm in this class is the mean shift algorithm, initially introduced by Fukunaga and Hostetler (1975) and subsequently recast by Cheng (1995) and Comaniciu and Meer (2002) to highlight several applications in Engineering. The mean shift algorithm is an iterative procedure which, at every step, moves the point obtained in the previous iteration to a location of higher density, thus producing a convergent sequence that transports any initial value to a local maximum of the density. Explicitly, any initial point  $\mathbf{y}_0$  is transformed recursively to obtain a sequence defined by

$$\mathbf{y}_{j+1} = \mathbf{y}_j + \mathbf{A} \mathbf{D}f(\mathbf{y}_j) / f(\mathbf{y}_j), \quad (3)$$

where  $\mathbf{A}$  is a positive-definite matrix chosen to ensure convergence. Notice that if  $\mathbf{A}$  is a positive multiple of the identity matrix then the sequence approximately follows the steepest ascent path, since at every step the shift is made along the gradient direction. In this case, the mean shift algorithm is easily recognized as a variant of the classical gradient ascent algorithm used for numerical maximization, but employing the normalized gradient (the gradient divided by  $f$ ) to accelerate convergence in low-density zones.

Since this algorithm can be applied to any initial point, it effectively produces a partition of  $\mathbb{R}^d$ . However, when a sample  $\mathbf{X}_1, \dots, \mathbf{X}_n$  is available it is common to use every sample point as initial value to obtain a clustering of the data. On the other hand, if the sample is employed to construct a smooth estimator  $\hat{f}$  of the density, then Equation (3) with  $f$  replaced by  $\hat{f}$  allows to produce a data-based clustering (hence, also a clustering of the data as noted above).

Moreover, for certain types of density estimators there exist closed forms for the matrix  $\mathbf{A}$  that ensure the convergence of the algorithm to a local maximum. In general, suppose the density  $f$  corresponds to a normal mixture model, so that

$$f(\mathbf{x}) = \sum_{\ell=1}^L \pi_{\ell} \phi(\mathbf{x} | \boldsymbol{\mu}_{\ell}, \boldsymbol{\Sigma}_{\ell}), \quad (4)$$



where  $\phi(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Sigma}) = |2\pi\boldsymbol{\Sigma}|^{-1/2} \exp\{-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\}$  refers to the density of the normal distribution with mean vector  $\boldsymbol{\mu}$  and covariance matrix  $\boldsymbol{\Sigma}$ , and  $\pi_1, \dots, \pi_L$  are nonnegative quantities that sum to one. Then, Carreira-Perpiñán (2007) showed that the iterative algorithm defined by (3) is convergent for every initial value  $\mathbf{y}_0$  by taking  $\mathbf{A} \equiv \mathbf{A}_j = \boldsymbol{\Sigma}_*(\mathbf{y}_j)$ , where  $\boldsymbol{\Sigma}_*(\mathbf{y}) = \{\sum_{\ell=1}^L \alpha_\ell(\mathbf{y}) \boldsymbol{\Sigma}_\ell^{-1}\}^{-1}$  is a weighted harmonic mean of the covariance matrices of the components, with weights  $\alpha_\ell(\mathbf{y}) = \pi_\ell \phi(\mathbf{y}|\boldsymbol{\mu}_\ell, \boldsymbol{\Sigma}_\ell)/f(\mathbf{y})$  given by the posterior odds of the  $\ell$ th component of the mixture after observing  $\mathbf{y}$ .

Hence, in the normal mixture model where the density estimator  $\hat{f}$  is obtained by replacing the parameters in (4) with estimates  $\hat{\pi}_\ell, \hat{\boldsymbol{\mu}}_\ell, \hat{\boldsymbol{\Sigma}}_\ell$ , the mean shift algorithm with  $\mathbf{A} = \hat{\boldsymbol{\Sigma}}_*(\mathbf{y}_j)$  (i.e., the previous definition, but also with the parameters replaced by estimates) provides a computationally simple way to obtain the corresponding data-based clustering. Some examples of the partitions thus obtained, for normal mixture densities taken from Wand and Jones (1993) and Ray and Lindsay (2005), are shown in Figure 7.

Furthermore, suppose that a nonparametric kernel density estimator  $\hat{f}_{\mathbf{H}}$  is employed in (3), namely  $\hat{f}_{\mathbf{H}}(\mathbf{x}) = n^{-1} \sum_{i=1}^n K_{\mathbf{H}}(\mathbf{x} - \mathbf{X}_i)$ , where the kernel  $K$  is a fixed density function,  $\mathbf{H}$  is a positive-definite bandwidth matrix and  $K_{\mathbf{H}}(\mathbf{x}) = |\mathbf{H}|^{-1/2} K(\mathbf{H}^{-1/2} \mathbf{x})$ . If the kernel  $K$  is taken as the normal density, then  $\hat{f}_{\mathbf{H}}(\mathbf{x}) = \sum_{i=1}^n n^{-1} \phi(\mathbf{x}|\mathbf{X}_i, \mathbf{H})$ , and so its structure as a normal mixture density allows to conclude as before that the mean shift algorithm is convergent for any initial point by just taking  $\mathbf{A} = \mathbf{H}$ . More generally, proceeding as in the proof of Theorem 1 in Comaniciu and Meer (2002) it can be showed that convergence is guaranteed with  $\mathbf{A} = \mathbf{H}$  for a general kernel of the form  $K(\mathbf{x}) = k(\mathbf{x}^\top \mathbf{x})$ , as long as the profile  $k : [0, \infty) \rightarrow \mathbb{R}$  is convex and monotonically decreasing.

Many variants of the mean shift algorithm have been explored in the literature, expanding in different directions. In the nonparametric context, the iterative procedure (3) with  $\mathbf{A} = \mathbf{H}$  as above can be explicitly written as  $\mathbf{y}_{j+1} = \sum_{i=1}^n W_i(\mathbf{y}_j) \mathbf{X}_i$ , hence taking the form of a weighted mean of the observations, with the weight functions defined as

$$W_i(\mathbf{y}) = \frac{G_{\mathbf{H}}(\mathbf{y} - \mathbf{X}_i)}{\sum_{i=1}^n G_{\mathbf{H}}(\mathbf{y} - \mathbf{X}_i)},$$

where the kernel  $G(\mathbf{x}) \propto g(\mathbf{x}^\top \mathbf{x})$  is based on the positive profile  $g = -k'$ . Sheik, Khan and Kanade (2007) provided an alternative formulation of (3) by noting that

$$\mathbf{y}_{j+1} = \underset{\mathbf{y} \in \mathbb{R}^d}{\operatorname{argmin}} \sum_{i=1}^n \|\mathbf{X}_i - \mathbf{y}\|_2^2 G_{\mathbf{H}}(\mathbf{X}_i - \mathbf{y}_j), \quad (5)$$

where  $\|\mathbf{x}\|_2 = (\mathbf{x}^\top \mathbf{x})^{1/2}$  denotes the usual Euclidean distance, and proposed the medioid-shift algorithm by restricting the minimization in (5) to the set of sample points, thus reducing the computation time since only one iteration of the algorithm has to be computed per data point because  $\mathbf{y}_{j+1} = \mathbf{X}_k$  for some  $k \in \{1, \dots, n\}$ . Accelerating the mean shift algorithm is in fact one of the expansion directions that has generated much research. A




Figure 7: Examples showing the ideal population clustering for several bivariate normal mixture densities, calculated using the mean shift algorithm.

(clearly non-exhaustive) list of references on this topic would include Carreira-Perpiñán (2006), Wang *et al.* (2007) and Vedaldi and Soatto (2008).

Other variants of the mean shift algorithm are oriented towards robustness, by considering median shifts instead of mean shifts. Wang, Qiu and Zamar (2007) proposed the algorithm CLUES that at each iteration shifts the previous point to the coordinate-wise (i.e., marginal) median of its  $K$ -nearest neighbors and Shapira, Avidan and Shamir (2009) explored other possibilities using the different concepts of multivariate median related to depth measures. A further refinement of CLUES, with improved computational efficiency, and also based on local medians is the ATTRACTORS algorithm of Peña, Viladomat and Zamar (2012). In fact, it would be possible to modify (5) to obtain shift algorithms based on other norm-based local medians, which would be the local analogues to the global proposals

studied in Dodge and Rousson (1999) (see also Small, 1990).

Indeed, a geometric median shift algorithm over Riemannian manifolds was introduced in Wang and Huang (2010), and its applications to clustering were explored in Wang, Huang and Wu (2013). In this respect, the adaptation of shift algorithms to manifolds does represent a third major avenue of active research on the topic. For instance, among some other recent papers dealing with this subject we could cite Subbarao and Meer (2009), Ozertem and Erdogmus (2011) and Genovese *et al.* (2012).

## 5 Conclusions and directions for further research

At the time of comparing different clustering procedures it is necessary to have a “ground truth”, or population goal, that represents the ideal clustering to which the clustering algorithms should try to get close. Sometimes this ideal population clustering is not so easy to specify and it depends on the notion of cluster in which the researcher is interested. Here, the population goal for modal clustering is accurately identified, making use of some tools from Morse theory, as the partition of the space induced by the domains of attraction of the local maxima of the density function.

This definition needs the probability density to be smooth to a certain degree, specifically it must be a Morse function. It would be appealing to extend this notion to density functions that are not Morse functions, meaning either that they are smooth but have degenerate critical points or even that they are not differentiable. To treat the first case, it might be useful to resort to the theory of singularities of differential mappings, which is exhaustively covered in the book by Arnold *et al.* (1998), for instance. On the other hand, the study of the non-smooth case might start from Agrachev, Pallaschke and Scholtes (1997), where Morse theory for piecewise smooth functions is presented. Here, the key role would be played by the subgradient, which generalizes the concept of gradient for non-smooth functions.

Another interesting challenge consists of studying the choice of the parameters for the density estimators (the bandwidth for kernel estimators, the mixture parameters for mixture model estimators) that minimize the distance between the corresponding data-based clustering and the true population clustering, as measured by any of the distance between clusterings discussed in Section 3. Or, perhaps even better, to develop methods aimed to perform modal clustering that do not necessarily rely on a pilot density estimate, as it happens for classification methods, which are not necessarily based on a regression estimate.

On the algorithmic side, the alternative formulation (5) of the mean shift algorithm by Sheik, Khan and Kanade (2007) surely deserves further attention. Its resemblance to locally weighted regression methodology (Ruppert and Wand, 1994), identifying somehow mean shift as a locally constant adjustment, suggests that perhaps a local linear mean shift could improve the performance of clustering based on kernel density estimation, in the same way as local linear regression improves over the Nadaraya-Watson (locally constant)

estimator.

**Acknowledgments.** This work has been partially supported by Spanish Ministerio de Ciencia y Tecnología project MTM2010-16660. The author wishes to thank Prof. Antonio Cuevas from Universidad Autónoma de Madrid as well as Prof. Ricardo Faro from Universidad de Extremadura for insightful conversations and suggestions concerning the material of Section 3. The paper by Ray and Lindsay (2005), in which interesting connections between Morse theory and the topography of multivariate normal mixtures are illustrated, was thought-provoking enough to inspire part of this paper.

## References

- Ackerman, M. and Ben-David, S. (2008) Measures of clustering quality: a working set of axioms for clustering. *Proceedings of the Neural Information Processing Systems Conference (NIPS 2008)*.
- Agrachev, A.A., Pallaschke, D. and Scholtes, S. (1997) On Morse theory for piecewise smooth functions. *Journal of Dynamical and Control Systems*, **3**, 449–469.
- Ankerst M., Breuning M.M., Kriegel H.P. and Sander J. (1999) OPTICS: ordering points to identify the clustering structure. In: *International Conference on Management of Data (SIGMOD99)*, ACM, pp. 49–60.
- Arabie, P. and Boorman, S.A. (1973) Multidimensional scaling of measures of distance between partitions. *Journal of Mathematical Psychology*, **10**, 148–203.
- Arnold, V.I., Goryunov, V.V., Lyashko, O.V. and Vasil’ev, V.A. (1998) *Singularity Theory I*. Springer-Verlag, Berlin.
- Azzalini, A. and Torelli, N. (2007) Clustering via nonparametric density estimation. *Stat. Comput.*, **17**, 71–80.
- Baíllo, A., Cuesta-Albertos, J.A. and Cuevas, A. (2001) Convergence rates in nonparametric estimation of level sets. *Statistics and Probability Letters*, **53**, 27–35.
- Baudry, J.-P., Raftery, A.E., Celeux, G., Lo, K. and Gottardo, R. (2010) Combining mixture components for clustering. *Journal of Computational and Graphical Statistics*, **19**, 332–353.
- Ben-David, S., von Luxburg, U. and Pál, D. (2006) A sober look at clustering stability. In G. Lugosi and H.-U. Simon, editors, *Proceedings of the 19th Annual Conference on Learning Theory (COLT)*, pp. 5–19. Springer, Berlin.

- Biernacki, C., Celeux, G. and Govaert, G. (2000) Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **22**, 719–725.
- Burkard, R., Dell’Amico, M. and Martello, S. (2009) *Assignment Problems*. SIAM, Philadelphia.
- Cadre, B. (2006) Kernel estimation of density level sets. *Journal of Multivariate Analysis*, **97**, 999–1023.
- Cadre, B., Pelletier, B. and Pudlo, P. (2012) Clustering by estimation of density level sets at a fixed probability. Preprint, available at <http://hal.archives-ouvertes.fr/hal-00397437>
- Carreira-Perpiñán, M.Á. (2006) Acceleration strategies for Gaussian mean-shift image segmentation. *Proceedings of the 2006 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1160–1167.
- Carreira-Perpiñán, M.Á. (2007) Gaussian mean-shift is an EM algorithm. *IEEE Transactions On Pattern Analysis And Machine Intelligence*, **29**, 767–776.
- Charon, I., Dencoud, L., Guénoche, A. and Hudry, O. (2006) Maximum transfer distance between partitions. *Journal of Classification*, **23**, 103–121.
- Cheng, Y. (1995) Mean shift, mode seeking, and clustering. *IEEE Transactions On Pattern Analysis And Machine Intelligence*, **17**, 790–799.
- Comaniciu, D. and Meer, P. (2002) Mean shift: A robust approach toward feature space analysis. *IEEE Transactions On Pattern Analysis And Machine Intelligence*, **24**, 603–619.
- Cuevas, A., Febrero, M. and Fraiman, R. (2001) Cluster analysis: a further approach based on density estimation. *Computational Statistics and Data Analysis*, **36**, 441–459.
- Cuevas, A. and Fraiman, R. (1997) A plug-in approach to support estimation. *Annals of Statistics*, **25**, 2300–2312.
- Cuevas, A. and Fraiman, R. (2010) Set estimation. In W. Kendall and I. Molchanov, editors, *New Perspectives in Stochastic Geometry*, pp. 374–397. Oxford University Press, Oxford.
- Day, W.H.E. (1981) The complexity of computing metric distances between partitions. *Mathematical Social Sciences*, **1**, 269–287.
- Dencoud, L. (2008) Transfer distance between partitions. *Advances in Data Analysis and Classification*, **2**, 279–294.

- Devroye, L., Györfi, L. and Lugosi, G. (1996) *A Probabilistic Theory of Pattern Recognition*. Springer-Verlag, New York.
- Dodge, Y. and Rousson, V. (1999) Multivariate  $L_1$ -mean. *Metrika*, **49**, 127–134.
- Ester M., Kriegel H.P., Sander J. and Xu X. (1996) A density-based algorithm for discovering clusters in large spatial databases with noise. In: *Proceedings of the 2nd International Conference on Knowledge Discovery in Data Mining (KDD-96)*, Portland, OR, USA. ACM, pp. 226–231.
- Everitt, B.S., Landau, S., Lesse, M. and Stahl, D. *Cluster Analysis* (fifth edition). John Wiley & Sons, Chichester.
- Fraley, C. and Raftery, A.E. (2002) Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association*, **97**, 611–631.
- Fukunaga, K. and Hostetler, L.D. (1975) The estimation of the gradient of a density function, with applications in pattern recognition. *IEEE T. Inform. Theory*, **21**, 32–40.
- García-Escudero, L.A., Gordaliza, A., Matrán, C. and Mayo, A. (2010) A review of robust clustering methods. *Advances in Data Analysis and Classification*, **4**, 89–109.
- Genovese, C.R., Perone-Pacifico, M., Verdinelli, I. and Wasserman, L. (2012) Nonparametric ridge estimation. *arXiv preprint arXiv:1212.5156*.
- Hand, D., Mannila, H. and Smyth, P. (2001) *Principles of Data Mining*. The MIT Press, Cambridge.
- Hartigan, J.A. (1975) *Clustering Algorithms*. Wiley, New York.
- Hartigan, J. A. (1987). Estimation of a convex density contour in two dimensions. *Journal of the American Statistical Association*, **82**, 267–270.
- Hastie, T., Tibshirani, R. and Friedman, J. (2009) *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (2nd edition). Springer-Verlag, New York.
- Hennig, C. (2010) Methods for merging Gaussian mixture components. *Advances in Data Analysis and Classification*, **4**, 3–34.
- Hyndman, R. (1996. Computing and graphing highest density regions. *The American Statistician*, **50**, 120–126.
- Izenman, A.J. (2008) *Modern Multivariate Statistical Techniques: Regression, Classification and Manifold Learning*. Springer Science+Business Media, New York.
- Jost, J. (2011) *Riemannian Geometry and Geometric Analysis*. Springer, Heidelberg.

- Li, J., Ray, S. and Lindsay, B.G. (2007) A nonparametric statistical approach to clustering via mode identification. *Journal of Machine Learning Research*, **8**, 1687–1723.
- MacQueen, J.B. (1967) Some methods for classification and analysis of multivariate observations. *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, **1**, 281–297. University of California Press, Berkeley.
- Mason, D. and Polonik, W. (2009) Asymptotic normality of plug-in level set estimates. *Annals of Applied Probability*, **19**, 1108–1142.
- Matsumoto, Y. (2002) *An Introduction to Morse Theory*. American Mathematical Society, Providence.
- Meilă, M. (2005) Comparing clusterings—an axiomatic view. In S. Wrobel and L. De Raedt, editors, *Proceedings of the International Machine Learning Conference (ICML)*. ACM Press, New York.
- Meilă, M. (2007) Comparing clusterings—an information based distance. *Journal of Multivariate Analysis*, **98**, 873 – 895.
- Meilă, M. (2012) Local equivalences of distances between clusterings—a geometric perspective. *Machine Learning*, **86**, 369–389.
- Milnor, J. (1963) *Morse Theory*. Princeton University Press, Princeton.
- Müller, D. W. and Sawitzki, G. (1987) Using excess mass estimates to investigate the modality of a distribution. Preprint 398, SFB 123, Universität Heidelberg.
- Müller, D. W. and Sawitzki, G. (1991) Excess mass estimates and tests for multimodality. *Journal of the American Statistical Association*, **86**, 738–746.
- Nugent, R. and Stuetzle, W. (2010) Clustering with confidence: a low-dimensional binning approach. In H. Locarek-Junge and C. Weihs, editors, *Classification as a Tool for Research*, pp. 117–125. Springer-Verlag, Berlin.
- Ozertem, U. and Erdogmus, D. (2011) Locally defined principal curves and surfaces. *Journal of Machine Learning Research*, **12**, 241–274.
- Peña, D., Viladomat, J. and Zamar, R. (2012) Nearest-neighbors medians clustering. *Statistical Analysis and Data Mining*, **5**, 349–362.
- Pollard, D. (1981) Strong consistency of  $k$ -means clustering. *Annals of Statistics*, **9**, 135–140.
- Polonik, W. (1995) Measuring mass concentrations and estimating density contour clusters—an excess mass approach. *Annals of Statistics*, **23**, 855–881.

- Ray, S. and Lindsay, B.G. (2005) The topography of multivariate normal mixtures. *Annals of Statistics*, **33**, 2042–2065.
- Rinaldo, A., Singh, A., Nugent, R. and Wasserman, L. (2012) Stability of density-based clustering. *Journal of Machine Learning Research*, **13**, 905–948.
- Rinaldo, A. and Wasserman, L. (2010) Generalized density clustering. *Annals of Statistics*, **38**, 2678–2722.
- Rodríguez-Casal, A. (2003) *Estimación de Conjuntos y sus Fronteras. Un Enfoque Geométrico*. Ph.D. thesis, University of Santiago de Compostela.
- Ruppert, D. and Wand, M.P. (1994) Multivariate locally weighted least squares regression. *Annals of Statistics*, **22**, 1346–1370.
- Samworth, R.J. and Wand, M.P. (2010) Asymptotics and optimal bandwidth selection for highest density region estimation. *Annals of Statistics*, **38**, 1767–1792.
- Shapira, L., Avidan, S. and Shamir, A. (2009) Mode-detection via median shift. *Proceedings of the 2009 IEEE International Conference on Computer Vision*, pp. 1–8.
- Sheikh, Y.A., Khan, E.A. and Kanade, T. (2007) Mode-seeking by medoidshifts. In *Proceedings of the 2007 IEEE International Conference on Computer Vision*.
- Small, C.G. (1990) A survey of multivariate medians. *International Statistical Review*, **58**, 263–277.
- Stuetzle, W. (2003) Estimating the cluster tree of a density by analyzing the minimal spanning tree of a sample. *Journal of Classification*, **20**, 25–47.
- Stuetzle, W. and Nugent, R. (2010) A generalized single linkage method for estimating the cluster tree of a density. *J. Comput. Graph. Statist.*, **19**, 397–418.
- Subbarao, R. and Meer, P. (2009) Nonlinear mean shift over Riemannian manifolds. *International Journal of Computer Vision*, **84**, 1–20.
- Thom, R. (1949) Sur une partition en cellules associée à une fonction sur une variété. *Comptes Rendus Hebdomadaires des Séances de l'Académie des Sciences*, **228**, 973–975.
- Tsybakov, A.B. (1997) On nonparametric estimation of density level sets. *Annals of Statistics*, **25**, 948–969.
- Vedaldi, A. and Soatto, S. (2008) Quick shift and kernel methods for mode seeking. In D. Forsyth, P. Torr and A. Zisserman, editors, *Computer Vision – ECCV 2008*, pp. 705–718. Springer-Verlag, Berlin.



- Vitalli, M. (2010) *Morse Decomposition of Geometric Meshes with Applications*. Ph. D. thesis, Università di Genova.
- von Luxburg, U. (2004) *Statistical Learning with Similarity and Dissimilarity Functions*. Ph.D. thesis, Technical University of Berlin.
- von Luxburg, U. and Ben-David, S. (2005) Towards a statistical theory for clustering. In *PASCAL workshop on Statistics and Optimization of Clustering*.
- Walther, G. (1997) Granulometric smoothing. *Annals of Statistics*, **25**, 2273–2299.
- Wand, M.P. and Jones, M.C. (1993) Comparison of smoothing parameterizations in bivariate kernel density estimation. *Journal of the American Statistical Association*, **88**, 520–528.
- Wand, M.P. and Jones, M.C. (1995) *Kernel Smoothing*. London, Chapman & Hall.
- Wang, Y. and Huang, X. (2010) Geometric median-shift over Riemannian manifolds. In B.-T. Zhang and M.A. Orgun, editors, *PRICAI 2010: Trends in Artificial Intelligence*, pp. 268–279. Springer-Verlag, Berlin.
- Wang, Y., Huang, X. and Wu, L. (2013) Clustering via geometric median shift over Riemannian manifolds. *Information Sciences*, **220**, 292–305.
- Wang, P., Lee, D., Gray, A. and Rehg, J.M. (2007) Fast mean shift with accurate and stable convergence. *JMLR Workshop and Conference Proceedings*, **2**, 604–611.
- Wang, X., Qiu, W. and Zamar, R.H. (2007) CLUES: A non-parametric clustering method based on local shrinking. *Computational Statistics and Data Analysis*, **52**, 286–298.
- Wishart, D. (1969) Mode analysis: a generalization of nearest neighbor which reduces chaining effects. In A.J. Cole, editor, *Numerical Taxonomy*, pp. 282–311. Academic Press, New York.
- Zadeh, R.B. and Ben-David, S. (2009) A uniqueness theorem for clustering. *UAI'09 Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, 639–646.